

REMARKS

Claims 8, 9, 11-13, and 19-25 are currently pending in the application. Claims 10 and 14-18 have been canceled. Claim 8 has been amended. Claims 19-25 have been added. Applicant respectfully submits that no new matter has been added. Applicant respectfully requests reconsideration of the application in view of the following remarks.

Claims 8-13 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Barthelemy ("Small Worlds Networks: Evidence for a Crossover Picture") in view of Collins ("It's a small world"). Claim 8 relates to a method for constructing a scalable computer system.

The Office Action asserts that interconnecting a plurality of computing nodes to form a plurality of node clusters (start from a regular network with n vertices [nodes] each connected to z neighbors, forming neighborhoods [node clusters])(p. 3181, col. 2); and "providing a plurality of cross-links between the node clusters (each node is connected to z neighbors, forming cross links between node clusters)(p. 3181, col. 2)." Applicant respectfully disagrees.

Barthelemy defines the word neighborhood as "...share with regular networks the concept of neighborhood, which means that if vertices i and j are neighbors then they will have many common neighbors..." (p. 3180, col. 1). Applicant, on the other hand, uses "neighborhood" and "cluster" in an ordinary sense, and states "[a] small world telecommunications network in accordance with the present invention is one in which vertices or nodes in local 'neighborhoods' or 'clusters' are interconnected to each other, and a relatively small number of random links or connections are provided between nodes of different neighborhoods of the network." In Barthelemy, the starting point is a regular network with n vertices, every one of which is connected to z neighbors. (p. 3181, col. 2) This is not equivalent to the previously claimed "interconnecting a plurality of computing nodes to form a plurality of node clusters."

Notwithstanding the foregoing remarks, claim 8 is amended to clarify aspects of the construction of a scalable computer system in accordance with the present invention. In particular, the method of construction comprises providing a first plurality of computing nodes, and a second plurality of switching nodes, wherein the plurality of computing nodes is more

numerous than the plurality of switching nodes. Each of the node clusters is then formed by interconnecting a corresponding group of the computing nodes via one of the switching nodes.

Accordingly, the starting point for the construction of a scalable computer system in accordance with the present invention is the formation of a plurality of disjoint node clusters, wherein a plurality of nodes are interconnected via a common switching node. This specific method of construction, and the scalable system produced thereby, is not disclosed or suggested in the prior art.

In paragraph 3 of the Detailed Action (“Response to Arguments”) the Examiner specifically refers to Applicant’s previous assertion that the clusters in the present invention are “interconnected using a necessarily smaller number of links”, noting that this limitation did not previously appear in the claims, and that Applicant has not clearly defined “node clusters” in the specification or the claims. Applicant submits that the limitations now introduced to clarify claim 8 address the Examiner’s concerns in this regard. In particular, it is now specifically recited that the node clusters are groups of computing nodes interconnected via corresponding common switching nodes. Furthermore, it is now clear that the number of switching nodes, corresponding with the number of node clusters, and hence the number of potential links between node clusters, is necessarily less numerous than the number of computing nodes. It is therefore believed that the amendments serve to clarify the features distinguishing the present invention from the cited prior art.

In the instant case, a person of ordinary skill in the art having common sense at the time of the invention would not have reasonably looked to the methods of Barthelemy to solve a problem in the construction of scalable systems. As the Office Action points out, Barthelemy and the Collins reference speculate that improvements in the operation and/or analysis of existing information networks may be achieved. They do not suggest the implementation of entirely new, scalable, computing systems. Neither Barthelemy nor Collins relate to, or mention, the specific technical field of scalable multiprocessing computer systems. The prior art relied upon in the Office Action arises from research in the fields of biology and sociology, and is primarily mathematical and abstract in nature. Applicant respectfully submits that speculative references to potential applications in fields beyond the expertise of the authors, and which do not include the application covered by the present claims, do not suffice to render the claimed

subject matter obvious. Additionally, there is no suggestion to combine the teachings and suggestions of Barthelemy and Collins as advanced in the Office Action, except from using Applicant's invention as a template through a hindsight reconstruction of Applicant's claims.

The Office Action cites to Barthelemy, at p. 3181, col. 2, as disclosing two separate steps of the claimed method. However, there is an additional level of structure in the network resulting from the claimed method, as compared with the network of Barthelemy, now clarified by the amendment of claim 8.

Claim 8 relates to a method for constructing a scalable computer system. Applicant respectfully submits that the cited combination of Barthelemy and Collins fails to teach, suggest, or render obvious at least one of the distinguishing features of independent claim 8, namely, providing a first plurality of computing nodes, and a second plurality of switching nodes, said first plurality being more numerous than said second plurality, forming a plurality of node clusters, each node cluster comprising one of said switching nodes interconnecting a corresponding group of said computing nodes, providing a plurality of cross-links between the node clusters, and directly connecting, using the plurality of cross-links, a corresponding plurality of pairs of node clusters selected from the plurality of node clusters in accordance with a selection process resulting in a formation of a network of the plurality of computing nodes having a higher clustering coefficient of nodes in comparison with a corresponding randomly-connected network in combination with a lower characteristic path length between the nodes. In addition, Applicant respectfully submits that the cited references fail to disclose wherein the steps of providing the plurality of cross-links and directly connecting the plurality of pairs of node clusters in accordance with the selection process are repeated until the resulting network comprises a small-world network having an average path length between the plurality of nodes falling within a predetermined desired range, independently of a number of said plurality of computing nodes.

Applicant respectfully submits that historically there have been two main approaches for interconnecting large number of processing nodes. First, a large number of direct connections may be provided between nodes, resulting in a low mean path length between the nodes. This approach results in a substantially quadratic increase in the number of interconnections as the number of nodes are scaled. Second, regular lattice or ring type

structures employing mainly local connections between nodes enable a manageable scaling of interconnections as the number of nodes increase. However, this approach requires messages to be routed along multi-hop paths. The path length increases approximately linearly with the size of the system. In the latter approach, propagation and buffering delays ultimately limit the system performance, such that further increase in the number of processing nodes fails to increase in processing capacity.

Applicant respectfully submits that, in accordance with the prior art, the trade off between interconnect complexity and path length exemplified by these two extremes is unavoidable. Therefore, it is not possible to achieve a low mean path length in combination with a low interconnect complexity in a large scale system. As a result, the realization that the small world principle as claimed provides a methodology for designing and constructing massively scaleable systems having precisely this advantage is counter-intuitive in the face of the pre-existing beliefs, and represents a genuine breakthrough in the art.

Applicant respectfully submits that the new solution for interconnecting processing nodes to form scalable systems, as claimed, is quite different from prior art approaches. In particular, the method as defined, for example, in claim 8 comprises three distinct stages in the construction of a scalable computer system. In the first stage, a plurality of computing nodes is interconnected via a (smaller) plurality of switches to form a plurality of node clusters, each of which comprises computing nodes interconnected by a common switch. These node clusters are initially unconnected. Accordingly, in a second stage, cross links are provided between the node clusters, such that the clusters become connected. Finally, in a third stage, the cross links are used to directly connect pairs of clusters, which are selected so as to form a network having the desired properties. The second and third stages of the process may be repeated as necessary in order to achieve this goal.

Applicant respectfully submits that the cited art, whether considered individually or in combination, fails to teach, suggest or render obvious the three-stage process as claimed. In particular, the step of first interconnecting nodes via switches to form clusters is not disclosed in any of the cited references, and nor is the step of interconnecting node clusters in order to form a small-world network of the nodes comprising the clusters. The additional level of structure provided to the network via the clustering step contributes substantially to the very high level of

scalability enabled, up to at least 256,000,000 processors. Applicant respectfully submits that this degree of scalability is unprecedented in the art.

Dependent claims 9 and 11-13 depend from and further restrict independent claim 8 in a patentable sense. Applicant respectfully submits that, for at least the reasons set forth above with respect to the rejection of independent claim 8, dependent claims 9 and 11-13 distinguish over the cited combination of Barthemely and Collins and are in condition for allowance.

The rejection of claims 14-18 under 35 U.S.C. 103(a) as being unpatentable over Barthemely ("Small Worlds Networks: Evidence for a Crossover Picture") in view of Collins ("It's a small world") further in view of Brewer (U.S. Pat. No. 5,859,975), is moot in view of the cancellation of these claims.

New independent claim 19 relates to a scalable computer system, incorporating limitations substantially corresponding with the method defined by claim 8. Applicant respectfully submits that that prior art fails to teach, suggest, or render obvious the combination of features recited in claim 19, and in corresponding dependent claims 20 to 25..

In view of the above amendment, Applicant believes the pending application is in condition for allowance.

Dated: February 7, 2008

Respectfully submitted,

Electronic Signature: /Jeffrey A. Tinker/

Jeffrey A. Tinker

Registration No.: 58,807

WINSTEAD PC

P.O. Box 50784

Dallas, Texas 75201

(214) 745-5400

Attorneys For Applicant